

AUTONEGOTIATION BETWEEN 1000BASE-X AND 1000BASE-T

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application No. 60/256,116, filed December 15, 2000, which is hereby incorporated by reference.

FIELD OF THE INVENTION

[0002] The present invention relates to networks, and more particularly to autonegotiation circuits within the physical layer of devices that are connected by an Ethernet network.

BACKGROUND OF THE INVENTION

[0003] Ethernet is a shared local area (LAN) networking technology that was developed in the early 1970s. The basic design includes a shared transmission medium such as coaxial cable and optical fiber. Since the communication medium is shared, nodes listen to make sure that the cable is not in use before transmitting data. Ethernet also defines protocols to handle collisions that occur when two nodes transmit data simultaneously.

[0004] Ethernet switches are Layer 2 devices that provide a switching matrix or fabric that can temporarily connect a port to any other port. For example, a computer A is connected to port A of a switch. A computer D is connected to port D of the switch. The switch connects the ports A and D and

provides a contentionless, full-bandwidth link to the computers A and D. Since the computers A and D are the only devices that are connected to the link, they are the only devices that compete for the link. The switch forwards frames or packets from one port to another – unlike a hub that forwards frames or packets to all other ports. This reduces traffic and increases security. A group of computers or nodes may also share a single port of the switch. Multilayer switches combine routing functions with switching. Other switches provide half and full duplex modes. The full duplex mode allows systems to establish connections and to send and receive data across a separate twisted pair or other media.

[0005] Gigabit Ethernet networks provide transmission speeds of 1000 Mb/s and include modes such as 1000Base-X and 1000Base-T. 1000Base-X modes include 1000Base-LX (IEEE 802.3z) and 1000Base SX (IEEE 802.3z). 1000Base-LX implements long-wavelength (1310 nm) laser transmissions with links up to 550 meters over multimode fiber optic cable and 3000 meters over single mode fiber optic cable. 1000Base SX implements short-wavelength (850 nm) laser transmissions over multimode fiber optic cable. 1000Base-T provides transmissions over four pairs of category 5 cable with a maximum distance of 100 meters per station to a switch or 205 meters end to end.

[0006] Referring now to Figure 1, a flexible Ethernet switch 10 allows selection of the type of physical media that is used. The switch 10 includes a plurality of ports 12-1, 12-2, ..., and 12-n. Some of the ports such as ports 12-1, 12-2, 12-3, and 12-4 are fixed media ports. Other ports such as ports 12-5 and

12-6 are configurable media ports. Gigabit interface connector (GBIC) modules 16-1 and 16-2 are connected to one of the configurable ports 12-5 and/or 12-6. Since both sides (20-1 and 22-1 and 20-2 and 22-2) of the GBIC modules 16-1 and 16-2 are 1000BASE-X, autonegotiation information can be passed freely to physical layers 24-1 and 24-2 of devices 26-1 and 26-2 that are connected thereto, respectively.

[0007] Referring now to Figure 2, for purposes of clarity reference numbers from Figure 1 have been used in Figure 2 to identify similar elements. A GBIC module 40 includes a physical layer with an integrated serializer/deserializer (SERDES). The GBIC module 40 connects 1000BASE-X media 20-1 and 20-2 from the switch 10 to 1000BASE-T media 42-1 and 42-2. The 1000Base-T media 42-1 and 42-2, in turn, is connected to physical layers 44-1 and 44-2 of devices 46-1 and 46-2, respectively. For example, see U.S. Patent Application Serial No. 09/501,556, filed February 9, 2000 and assigned to the assignee of the present invention, which is hereby incorporated by reference. Since there is no direct path between 1000BASE-X and 1000BASE-T networks for passing autonegotiation information, 1000BASE-X autonegotiation is disabled.

[0008] 1000BASE-X autonegotiation information is exchanged by using special code groups that are not used during normal packet transmission. In 1000BASE-X autonegotiation, two devices (P and Q) communicate with each other over a link. The device Q is the link partner of the device P and the device P is the link partner of the device Q. 1000BASE-X autonegotiation uses the

underlying media to pass 16 bits of autonegotiation information at a time. The 16 bits of autonegotiation information are embedded in configuration ordered sets that are not used during normal data transmission. Therefore, devices are able to distinguish whether the transmitted data is a normal packet or autonegotiation data.

[0009] When the device P transmits a first configuration ordered set with all 16 bits set equal to zeros, the link partner Q knows that the device P is restarting autonegotiation. The device P continues to transmit the first configuration ordered set with all zeros until it is ready to start autonegotiation. Once the device P is ready to start autonegotiation, it transmits a second configuration ordered set with 16 bits of autonegotiation data. At this point, the second configuration ordered set is not all zeros. The device P continues to transmit the second configuration ordered set until the link partner Q transmits a second configuration ordered set that is not all zeros. When both of the devices P and Q are transmitting the second configuration ordered sets, autonegotiation continues according to the IEEE protocol for exchanging data.

[0010] There are currently 6 bits of the second configuration ordered set that are defined for 1000BASE-X media. These bits indicate 1000BASE-X full duplex, 1000BASE-X half duplex, pause, asymmetric pause, remote fault 1 and remote fault 2. A device advertises these capabilities by either setting or clearing the defined bits. The device may advertise that it has the function only if it can actually perform that function. However, the functionally capable device may choose to not advertise the function.

SUMMARY OF THE INVENTION

[0011] An autonegotiation circuit for Ethernet networks according to the invention includes a first device that is connected to a first media. A second device is connected to a second media. A network interface connector (NIC) module is connected to the first device by the first media and to the second device by the second media. The NIC module provides autonegotiation between the first and second devices.

[0012] In other features of the invention, the first device includes a first NIC interface including a transmitter and a receiver. The NIC module includes a second NIC interface with a transmitter and a receiver and a first copper interface with a transmitter and a receiver. The second device includes a second copper interface with a transmitter and a receiver. The transmitter of the first NIC interface is connected to the receiver of the second NIC interface. The receiver of the first NIC interface is connected to the transmitter of the second NIC interface. The transmitter of the first copper interface is connected to the receiver of the second copper interface. The receiver of the first copper interface is connected to the transmitter of the second copper interface.

[0013] In still other features, the transmitters of the first and second NIC interfaces transmit a first configuration ordered set. The receiver of the second NIC interface receives a second configuration ordered set from the transmitter of the first NIC interface. The NIC module stores in memory first configuration data of the first device that is contained in the second configuration ordered set. The transmitter of the first copper interface transmits a first FLP

burst. The first FLP burst contains at least one configuration parameter provided by the first configuration data.

[0014] In still other features, the first copper interface and the second copper interface complete autonegotiation by exchanging additional data and establishing a link. The transmitter of the second copper interface transmits second configuration data that is stored in the memory of the NIC module. The transmitter of the second NIC interface generates a second configuration ordered set that contains at least one configuration parameter provided by the second configuration data. The first and second NIC interfaces establish a link.

[0015] Further areas of applicability of the present invention will become apparent from the detailed description provided hereinafter. It should be understood that the detailed description and specific examples, while indicating the preferred embodiment of the invention, are intended for purposes of illustration only and are not intended to limit the scope of the invention.

BRIEF DESCRIPTION OF THE DRAWINGS

[0016] The present invention will become more fully understood from the detailed description and the accompanying drawings, wherein:

[0017] Figure 1 illustrates the interface between an Ethernet switch having a GBIC module and a device according to the prior art;

[0018] Figure 2 illustrates the interface between an Ethernet switch having a GBIC module with flexible media capability and a device according to the prior art;

[0019] Figure 3 illustrates the interface between an Ethernet switch having a GBIC module with flexible media capability and a device according to the present invention;

[0020] Figure 4 illustrates the interface between an Ethernet switch having a GBIC module with mixed media capability and a device according to the present invention;

[0022] Figures 5A and 5B are datagrams illustrating the exchange of autonegotiation information between a 1000BASE-X interface and a 1000BASE-T interface in accordance with the present invention; and

[0023] Figure 6 is a flowchart illustrating steps performed by the GBIC module during autonegotiation.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[0024] The following description of the preferred embodiment(s) is merely exemplary in nature and is in no way intended to limit the invention, its application, or uses.

[0025] Autonegotiation in both 1000BASE-T media and 1000BASE-X media includes 16 bits of autonegotiation data. In 1000Base-T media, however, the underlying mechanism of transmitting the bits is more complex. The 16 bits of data are embedded in a fast link pulse (FLP) burst. Generally, it takes more time to complete the 1000BASE-T autonegotiation than it does to complete the 1000BASE-X autonegotiation. In addition, there are many more bits defined for 1000BASE-T autonegotiation. In the GBIC application, only

1000BASE-T full duplex, 1000BASE-T half duplex, pause, and asymmetric pause functions are relevant. Other autonegotiation information that is required to establish a 1000BASE-T link (for example, master/slave information) has no relevance to 1000BASE-X. To restart autonegotiation on the 1000BASE-T side, the device wishing to restart breaks the communication link for over 1 second. During autonegotiation, the device does not transmit anything including fast link pulse bursts for over 1 second.

[0026] Referring now to Figure 3, a switch 50 has a GBIC slot 52. A GBIC module 56 plugs into the GBIC slot 52. The switch 50 and the GBIC module 56 communicate via the 1000BASE-X protocol through an interface G_s 58 (where G refers to GBIC and s refers to the switch 50), and an interface G_m 80 (where G refers to GBIC and m refers to the GBIC module 56). A device 62 is connected to the switch 50. The device 62 has a 1000BASE-T interface C_d 64 (where C refers to copper and d refers to the device 62). The GBIC module 56 also has a 1000BASE-T interface C_m 66 (where C refers to copper and m refers to the GBIC module 56). An example of the GBIC module 56 is the PCM-8519-3 that is available from Finisar and that has the 1000BASE-X interface G_m and the 1000BASE-T interface C_m . Currently, it is not possible for the switch 50 and the device 62 to exchange information about full duplex, half duplex or pause functions. The present invention solves this and other problems as will be described below.

[0027] Referring now to Figure 4, the interfaces G_s 58 and G_m 60 between the switch 50 and the GBIC module 56 and the interfaces C_m 66 and C_d

64 between the GBIC module 56 and the device 62 are shown. Each interface includes a transmitter and a receiver. The GBIC module 56 is preferably powered by the switch 50. Autonegotiation on the 1000BASE-X interface (G_s , G_m) is faster than the 1000BASE-T interface (C_m , C_d). Therefore, the GBIC module 56 will receive autonegotiation information faster from the switch 50 as compared to device 62.

[0028] Referring now to Figures 5A and 5B, information is passed between the 1000BASE-X interface and 1000BASE-T interface to allow autonegotiation. For ease of reference, section 1, section 2, ..., section n are labeled as they are referenced. In section 1, after power up both the interfaces G_s 58 and G_m 60 transmit first configuration ordered sets with all 0 data to restart the 1000BASE-X autonegotiation process. The interface C_d 64 sends fast link pulse (FLP) bursts to exchange 1000BASE-T autonegotiation information. However, unlike normal 1000BASE-T autonegotiation, the interface C_m 66 is disabled or in a wait mode until the interface G_s 58 transmits the first configuration ordered set with all zero data. Note that it may be possible that the interface C_d 64 is not sending anything at this time.

[0029] In section 2, the interface G_s 58 starts to transmit the second configuration ordered set with information on duplex and pause capabilities (in other words, not all bits are zero). However, the interface G_m 60 is still forced to transmit all zeros which forces the interface G_s 58 to wait. Note that it may be possible that the interface C_d 64 is not sending anything at this time.

[0030] In section 3, after the interface G_s 58 completes transmitting three consecutive second configuration ordered sets with identical data bits, the duplex and pause ability (and other functions that will be enabled in the future) of the switch 50 is stored by the GBIC module 56 in memory and these capabilities are advertised by the interface C_m 66. The interface C_m 66 starts transmitting (FLP) bursts with data advertising the capabilities of switch 50. Note that it may be possible that the interface C_d 64 is not sending anything at this time.

[0031] In section 4, the interfaces C_m 66 and C_d 64 complete the 1000BASE-T autonegotiation. There must be additional information exchanged beyond advertising the capabilities of the switch 50 to bring the 1000BASE-T link up. This additional information exchange is transparent to the switch 50. The device 62 sees the advertised capabilities of the switch 50.

[0032] In section 5, the GBIC module 56 stores the capabilities advertised by the interface C_d 64. The interface G_m 60 subsequently starts transmitting configuration ordered sets that advertise the capabilities of the device 62. The switch 50 receives the advertised capabilities of device 62. In section 6, the 1000BASE-X autonegotiation completes and the 1000BASE-X link is up. The link between the switch 50 and the device 62 is established and normal packet data can be transmitted and received.

[0033] If for any reason the GBIC module 56 detects that the 1000BASE-X link is down (for example, the switch 50 transmits the first configuration ordered set with all zeros), the GBIC module 56 stops transmitting on the interface C_m 66 for over a second. This causes the 1000BASE-T

autonegotiation to restart. The steps outlined above are executed again. If the GBIC module 56 detects that the 1000BASE-T link is down (for example, the device 62 stops transmitting), the GBIC module 56 transmits the first configuration ordered set with all zeros. This restarts the 1000BASE-X autonegotiation. The steps outlined above are executed again.

[0034] Referring now to Figure 6, steps performed by the GBIC module 56 are shown generally at 100. Control begins with step 102. In step 104, the interface G_m 60 transmits the first configuration ordered sets using the transmitter 76. In step 106, when the receiver 78 of the interface G_m receives the second configuration ordered set from the transmitter 72 of the interface G_s 58, the GBIC module 56 stores first configuration data in memory. In step 110, the transmitter 80 of the interface C_m 66 transmits FLP bursts with one or more parameters that are set based on the first configuration data. In step 112, if the receiver 82 of the interface C_m 66 receives a FLP burst from the device 62, control continues with step 114. Otherwise control loops to step 110.

[0035] In step 114, the interface C_m 66 completes autonegotiation with the interface C_d 64 and establishes a first link (the 1000BASE-T link). The transmitter 84 sends second configuration data to the GBIC module 56, which stores the second configuration data in memory. In step 116, the transmitter 76 of the interface G_m 60 transmits second configuration ordered sets that advertise the capabilities of the device 62. In step 118, the interface G_m 60 completes a second link (the 1000BASE-X link). If the copper link goes down, autonegotiation

is started on the fiber side. If the fiber link goes down, autonegotiation is started on the copper side.

[0036] Thus it will be appreciated from the foregoing that as a result of the present invention, an autonegotiation circuit and method for autonegotiation using first and second media is provided by which the principal objectives, among others, are completely fulfilled. It will be equally apparent and is contemplated that modification and/or changes may be made in the illustrated embodiment without departure from the invention. Accordingly, it is expressly intended that the foregoing description and accompanying drawings are illustrative of preferred embodiments only, not limiting, and that the true spirit and scope of the present invention will be determined by reference to the appended claims and their legal equivalent.